arXiv:1608.07048v1 [stat.CO] 25 Aug 2016

On the application of higher order symplectic integrators in Hamiltonian Monte Carlo

Janne Mannseth, Tore S. Kleppe, Hans J. Skaug

August 26, 2016

Abstract

We explore the construction of new symplectic numerical integration schemes to be used in Hamiltonian Monte Carlo and study their efficiency. Two integration schemes from Blanes et al. (2014), and a new scheme based on optimal acceptance probability, are considered as candidates to the commonly used leapfrog method. All integration schemes are tested within the framework of the No-U-Turn sampler (NUTS), both for a logistic regression model and a student t-model. The results show that the leapfrog method is inferior to all the new methods both in terms of asymptotic expected acceptance probability for a model problem and the and efficient sample size per computing time for the realistic models.

Keywords: Hamiltonian Monte Carlo; Leapfrog method; Parametrized numerical scheme; Integration scheme; No-U-Turn Sampler; Efficient Sample Size

1. Introduction

Hamiltonian Monte Carlo (Duane et al., 1987; Neal, 2010) (HMC) is a Markov chain Monte Carlo (MCMC) algorithm that avoids the random walk behavior of the commonly used Metropolis-Hastings algorithm. For this reason HMC has become popular for sampling from the posterior distribution of Bayesian models (Liu, 2008; Hoffman and Gelman, 2014).

HMC is applicable when the parameters have continuous distributions, and requires that the gradient of the log-posterior can be evaluated. Additional practical challenges

with the HMC algorithm is its sensitivity to two user specified parameters; the so-called step size and the number of steps. The NUTS algorithm (Hoffman and Gelman, 2014) automates the choice of these tuning parameters. There exist extensions of HMC, such as Riemann Manifold HMC and Lagrangian HMC (see e.g. Girolami and Calderhead (2011); Lan et al. (2015)). Here we employ plain HMC because of its widespread use e.g. via the software package STAN (Carpenter et al., 2015).

An essential component of HMC is to evolve the Hamiltonian system derived from the posterior distribution while sampling is taking place. In practice the pair of differential equations that describes this system must be discretized, while retaining the the volume-preserving and time-reverible nature of the exact solution (Neal, 2010). The most commonly used parametrized integration scheme used for HMC is the leapfrog method. Blanes et al. (2014) have developed a framework for deriving more efficient integration schemes. Such schemes have the potential to perform better than the leapfrog method, particularly in high dimensions. The leapfrog method often requires a small step size to obtian reasonable acceptance probabilites, whereas the new integration schemes have the potentential to allow larger step sizes while retianing good acceptance probabilites (Blanes et al., 2014).

The purpose of this paper is to compare the numerical integrators of Blanes et al. (2014), and also a new numerical integrator, to the leapfrog method by considering realistic problems and experiments. The effective sample size per computing time is used as the performance measure. The results indicate that the new integrators perform better than the leapfrog method, with neglible additional implementation effort.

The paper is laid out as follows: Section 2 defines HMC and in Section 3 the leapfrog method and the three other parameterized integration schemes are introduced. An asymptotic expression for the expected acceptance probability is found for each of the numerical schemes based on a Gaussian test problem. Section 4 embeds the integration schemes within the NUTS algorithm, and measures performance on a logistic regression model and a student t-model. Lastly, Section 5 contains some discussion.

2. Hamiltonian Monte Carlo

Denote the target (posterior) density by $\pi(q)$, where $q \in \mathbb{R}^d$. The system is augmented with another random vector $p \in \mathbb{R}^d$. In the language of Hamiltonian dynamics $q(t)$ and $p(t)$ represent respectively the position and momentum of a particle at time t. The Hamiltonian is denoted $H(q,p)$ and can be expressed as (Neal, 2010):

$$H(q,p) = U(q) + \frac{1}{2}p^T p, \tag{2.1}$$

where $U(q) = -\log(\pi(q))$. The last part of equation (2.1) represents the negative log density of a $\mathcal{N}(0,I)$ distribution where the density is defined as $f(\mathbf{x}) = (\frac{1}{\sqrt{2\pi}})^{\frac{n}{2}} \exp(-\frac{1}{2}x^T x)$ and I is the identity matrix. Note that strictly speaking only a kernel density proportional to the density is needed, not the density itself. The Hamiltonian governs the evaluation of $p(t)$ and $q(t)$ via the following system of differential equations:

$$\frac{dq}{dt} = \frac{\partial H}{\partial p} = p, \quad \frac{dp}{dt} = -\frac{\partial H}{\partial q} = -\nabla U(q), \tag{2.2}$$

where $\nabla U(q)$ is the gradient of $U(q)$.

Let (q,p) denote the current state of the sampler and (q^*,p^*) a proposal for the next. One iteration of the HMC algorithm involves the following steps. We start by drawing a new $p \sim \mathcal{N}(0,I)$. Next, we solve equation (2.2) using L steps of the integration scheme (leapfrog or a similar time-reversible, volume-preserving method). After L steps we arrive at the proposed position and momentum, q^* and p^*. The proposal is accepted (Duane et al., 1987; Neal, 2010) with probability:

$$\alpha = \min(1, \exp(\Delta)), \tag{2.3}$$

where $\Delta = (-H(q^*,p^*) + H(q,p))$. At each iteration p^* is discarded after α has been computed.

3. Symplectic numerical integration schemes

As described above, volume-preserving and time-reversible integrators are essential to HMC. Time is discretized in the differential equations (2.2) by introducing the step size ϵ. Each parameterized numerical scheme consists of partial steps updating either position or momentum. For each step it is indicated where we are in the process of going from time t to $t + \epsilon$.

3.1. The leapfrog method

Consider the first integration scheme known as the leapfrog method (Leimkuhler and Reich, 2004; Neal, 2010), expressed for notational simplicity for $t = 0$:

$$\begin{aligned} p(\epsilon/2) &= p(0) - \frac{\epsilon}{2}\nabla U(q(0)), \\ q(\epsilon) &= q(0) + \epsilon p(\epsilon/2), \\ p(\epsilon) &= p(\epsilon/2) - \frac{\epsilon}{2}\nabla U(q(\epsilon)). \end{aligned} \tag{3.1}$$

Each partial step represents an update in either p or q, jointly resulting in an update for (q, p) at time ϵ. To obtain a proposal, (3.1) is repeated L times, taking the output of the previous iteration as input to the next. The leapfrog method is widely used in practice, and therefore constitutes a reasonable reference.

3.2. Two-stage methods

Following the method suggested in Blanes et al. (2014) the next parametrized integration scheme, composed of two leapfrog steps, is calculated to be as follows:

$$\begin{aligned} q(a_1\epsilon) &= q(0) + a_1\epsilon p(0), \\ p(\epsilon/2) &= p(0) - \frac{\epsilon}{2}\nabla U(q(a_1\epsilon)), \\ q((1-a_1)\epsilon) &= q(a_1\epsilon) - (1-2a_1)\epsilon p(\epsilon/2), \\ p(\epsilon) &= p(\epsilon/2) - \frac{\epsilon}{2}\nabla U(q((1-a_1)\epsilon)), \\ q(\epsilon) &= q((1-a_1)\epsilon) + a_1\epsilon p(\epsilon). \end{aligned} \tag{3.2}$$

Following Blanes et al. (2014), the value of the parameter is $a_1 = \frac{3-\sqrt{3}}{6}$ (independent of the dimension d). This value is found by minimizing an expression for the energy error $H(q,p) - H(q^*,p^*)$ with respect to a_1. Inserting this a_1 before performing any further calculations gives the integration scheme called the "two-stage" method.

The next scheme has the same basis (3.2). We suggest an alternative criterion for determining the value for the parameter a_1. Now, a_1 is chosen so that the value of the expected acceptance probability α in (2.3) is maximized. As shown in Appendix A the optimal value is $a_1 = \frac{3-\sqrt{5}}{4}$, slightly smaller than for the two-stage method. Inserting this a_1 into equation (3.2) we get a new numerical scheme and name it the "new two-stage" method. The proposed method for finding the new value for a_1 is under the constraint that proposals are independent of the current state under a standard d-dimensional Gaussian model.

3.3. The three-stage method

The final integration scheme is called the three-stage method:

$$
\begin{aligned}
q(a_1\epsilon) &= q(0) + a_1\epsilon p(0),\\
p(b_1\epsilon) &= p(0) - b_1\epsilon\nabla U(q(a_1\epsilon)),\\
q(\epsilon/2) &= q(a_1\epsilon) + \left(\frac{1}{2} - a_1\right)\epsilon p(b_1\epsilon),\\
p((1-b_1)\epsilon) &= p(b_1\epsilon) - (1-2b_1)\epsilon\ \nabla U(q(\epsilon/2)),\\
q((1-a_1)\epsilon) &= q(\epsilon/2) + \left(\frac{1}{2} - a_1\right)\epsilon p((1-b_1)\epsilon),\\
p(\epsilon) &= p((1-b_1)\epsilon) - b_1\epsilon\nabla U(q((1-a_1)\epsilon)),\\
q(\epsilon) &= q((1-a_1)\epsilon) + a_1\epsilon p(\epsilon),
\end{aligned}
\tag{3.3}
$$

where the values of parameters a_1 and b_1 are found in Blanes et al. (2014) to be $a_1 = \frac{12127897}{102017882}$ and $b_1 = \frac{4271554}{14421423}$ using the energy error criterion. All the introduces integration schemes are easy to implement as alternatives to the leapfrog method.

3.4. Asymptotic acceptance probabilities

Based on an d-dimensional standard Gaussian test problem, we compute asymptotic (in dimension) expressions for the expected acceptance probability when the integration times, ϵL, correspond to proposals that are independent of the current state.

[t]Table 1 near here [/t]

Note that $\mathrm{E}(\alpha)$ is $1+O(\epsilon^4)$ for the new two-stage while it is $1+O(\epsilon^2)$ for the other three. Hence, the new two-stage method has the highest acceptance rate for sufficiently small ϵ, owing to how the value of parameter a_1 was found.

Beskos et al. (2013) provide a general result that shows that ϵ should behave as $O(d^{-\frac{1}{4}})$. To obtain $\mathrm{E}(\Delta) = O(1)$ for $L\epsilon = O(1)$, the results in Table 1 also indicate $\epsilon = O(d^{-\frac{1}{4}})$ for leapfrog, two-stage and three-stage under the standard Gaussian model problem. If ϵ decreases faster we will waste computation, and if it decreases slower, the MCMC chain will stagnate and be ineffective. On the other hand, the acceptance probability for the new two-stage method imply $\epsilon = O(d^{-\frac{1}{8}})$ to obtain $\mathrm{E}(\Delta) = O(1)$ for $L\epsilon = O(1)$. This could indicate that the new two-stage numerical scheme is too specialized to the particular model problem used to find it, and will be furhter explored in section 4. More details on the calculations in Table 1 are given in Appendix A.

3.5. Comparing performance on the model problem

To further add understanding of the results provided in Table 1, we compare a weighted measure of efficiency based on $\mathrm{E}(\alpha)$ for different numbers of dimensions. Again, we use a d-dimensional standard Gaussian model problem. Let Υ denote the expected number of accepted movements per calculation time, as a measure of efficiency. For the leapfrog method we have $\Upsilon = \frac{\mathrm{E}(\alpha)}{L}$, for the two-stage methods $\Upsilon = \frac{\mathrm{E}(\alpha)}{2L}$, where the factor 2 in the denominator is an effect of increase in number of gradient evaluations. Correspondingly the three-stage method has $\Upsilon = \frac{\mathrm{E}(\alpha)}{3L}$. The measure is considered for different values of dimension d and the maximum Υ within each dimension is the preferable one. For each dimension the

$$\max_{\epsilon} \Upsilon(\epsilon, d), \tag{3.4}$$

where ϵ is such that $L\epsilon \approx \frac{\pi}{2}$ corresponds to independent proposals, is found for all four numerical schemes. For each of the numerical schemes, Figure 1 shows how the maximum Υ develops as d increases. We see that the leapfrog method starts out the best, but is quickly overtaken by all the other schemes as d increases. It is also seen that the performance of the new two-stage method, which is of a higher order than the other schemes, stands out for high d. Figure 2 shows the ratio between Υ for the leapfrog method and each of the other numerical schemes. For each fraction, a value greater than one indicates that the suggested method (for the corresponding d) is the best of the two methods considered. In Figure 2 we see a different illustration on how the new methods perform better than leapfrog method as d increases. It is clear that the new two-stage method is the most efficient of all methods considered here .

[f]Figure 1 near here [/f]

[f]Figure 2 near here [/f]

Based on these results the way forward is to determine whether this is relevant in practice. Thus we need to consider some realistic examples.

4. Simulation studies and realistic applications

Arguably, the d-dimensional standard Gaussian problem (which Table 1 relies on) is too simple for us to claim the new integrators constitute improvements over leapfrog in practice. This section considers realistic models. In the d-dimensional Gaussian case, we used proposals independent of current state. In a general case that does not lend it self to analytic solution of (2.2), we resort to the No-U-Turn sampler for choosing appropriate integration times.

4.1. The No-U-Turn Sampler

The No-U-Turn Sampler (NUTS) was introduced by Hoffman and Gelman (2014) as an improvement to the HMC algorithm. One reason for using NUTS here is to avoid the sensitivity to user specified parameters, because both too low and too high values of step size ϵ and number of steps L can cause problems (Hoffman and Gelman, 2014). The NUTS algorithm discovers the point where the HMC algorithm could start to explore

random walk behaviour by allowing no u-turns in the trajectory. This way the choice of L is eliminated. In addition, a method for adapting the step size ϵ by using dual averaging is implemented in NUTS (Hoffman and Gelman, 2014). Using NUTS with adaptive selection of ϵ eliminates all user specified choices from Section 3. In the case of the four numerical schemes, this means that their efficiency can be tested without user specified parameters. All subsequent computations are carried out in MATLAB.

4.2. Simulation testing using NUTS and logistic regression model

As the first test case, we use a Bayesian logistic regression model, and the collection of five data sets and a Bayesian logistic regression model from Girolami and Calderhead (2011). The five data sets vary in number of observations n from 250 to 1000 and in number of parameters d from 7 to 25. The model layout is identical to Girolami and Calderhead (2011) where the prior is given as $\boldsymbol{\beta} \sim \mathcal{N}(\mathbf{0}, 100\mathbf{I})$. The effective sample size (ESS) is found by following Girolami and Calderhead (2011) and used in the performance measurement ESS/CPU time. All the experiments are repeated ten times, so that the results in Table 2 are mean results over these ten runs. Table 2 presents the mean CPU times, the mean minimum, median and maximum ESS and the mean number of step sizes for all numerical integration schemes and data sets. The final two columns of Table 2 show minimum ESS/CPU time and median ESS/CPU time. The number of samples were 5000 with burn-in set to 1000 for all experiments.

[t]Table 2 near here [/t]

In general Table 2 shows that the higher order numerical schemes get higher ESS/CPU time than the leapfrog method. The two-stage method generally performs better than the new two-stage method. Four data sets have ESS that reach the number of samples (5000). This means that there is little or no autocorrelation between the samples. Table 2 shows that the three-stage method or the two-stage method perform the best within each data set, as they have the largest minimum ESS/CPU time. The leapfrog method has the lowest value of minimum ESS/CPU time for all five data sets. The last column of Table 2, the median ESS/CPU time, confirms the same pattern as for the minimum ESS/CPU time. It is expected that the step size increases in relation to the number of steps the methods have. With regards to step size we have that the two-stage method

$>$ (the leapfrog method)/2 for all data sets except "Ripley". The same goes for the new two-stage method. Correspondingly we have that the three-stage method $>$ (the leapfrog method)/3 with regards to step size. For the "Ripley" data set this only holds for the new two-stage method, while the other two are very close. With these results we see that the suggested integration schemes outperform the leapfrog method also when considering a more realistic problem than the standard Gaussian test problem.

4.3. Simulation testing using NUTS and student t-model

In addition to the logistic model in Section 4.2, we also consider a multivariate student t distribution. The target density kernel is given as

$$\pi(x) \propto \left(1 + \frac{1}{\nu} x^T \Sigma^{-1} x\right)^{-\frac{\nu+d}{2}},$$

where the degrees of freedom, ν, is set to 5 and dimension d is $2, 10$ and 100. Σ^{-1} is the precision matrix associated with a Gaussian AR(1) model with autocorrelation 0.95 and unit innovation variance. All variations are run ten times and the results in Table 3 are the means taken over these ten runs.

[t]Table 3 near here [/t]

From Table 3 we see that in general the two-stage method performs better than the new two-stage method. For each dimension the highest minimum and median ESS is obtained with the two-stage method or the three-stage method. However, because of the time cost for the two-stage methods, the leapfrog method gets a slightly higher value of $\frac{\text{min ESS}}{\text{CPU time}}$ and $\frac{\text{med ESS}}{\text{CPU time}}$. The three-stage method requires less CPU time and gives the highest $\frac{\text{min ESS}}{\text{CPU time}}$ and $\frac{\text{med ESS}}{\text{CPU time}}$ of all four numerical integrators for dimensions $d = 10$ and $d = 100$. We see that for $d = 100$ that the step sizes of the two-stage methods and the three-stage method are greater than the step size of $\frac{\text{the leapfrog method}}{2}$ and $\frac{\text{the leapfrog method}}{3}$ respectively. For $d = 10$ this only holds for the three-stage method. For $d = 2$ it holds for the two-stage method. Considering that the integration schemes are all easy to implement and work with, these results give reason to further explore the use of them.

5. Discussion

This paper makes the contribution of introducing and developing numerical schemes as alternatives to the leapfrog method. For a d-dimensional standard Gaussian model all three new numerical integration schemes perform better than the original scheme as dimension d increases. It is seen that a numerical integration scheme of a higher order, like the new two-stage method, stands out for the Gaussian model, especially for larger d. Using these schemes can require more costly computations because of the number of gradient evaluations. However, they are still superior to their precursor in HMC.

For the logistic regression model the new schemes were all shown to be more efficient than the leapfrog method. Even with the increased gradient evaluations the results were positive regarding the efficiency of the suggested numerical integration schemes, for all the five datasets considered.

For the student t-model the suggested numerical integration schemes all obtain higher minimum and median ESSes than the leapfrog method for $d = 2$ and $d = 10$. For $d = 100$, the ESSes are quite close, but only the three-stage method has a higher minimum ESS than the leapfrog method. The leapfrog method has a higher minimum and median ESS/CPU time than the two-stage methods. This can be because of the cost related to the gradient evaluations of the two-stage methods, which can impact the CPU time. Still, the tree-stage method is the most efficient of all schemes as d increases. This means that also for the student t-model there can be efficiency gain by replacing the leapfrog method with other numerical schemes.

These results give reason to explore the development of numerical schemes on the form seen in this paper. By doing this the performance and efficiency of HMC can be improved, especially as the dimension increases.

Appendix A. Details

The $U(q)$ used when obtaining the results in Section 3 is $U(q) = \frac{1}{2}q^T q$ and the Hamiltonian then is:

$$H = \frac{1}{2}q^T q + \frac{1}{2}p^T p,$$

where $q, p \in \mathbb{R}^d$. Owing to the fact that the gradient of H with respect to both q and p are linear, H together with a numerical integration scheme gives a 2×2 matrix M_ϵ (one for each scheme) whose elements change with step size ϵ so that

$$\begin{bmatrix} q_i(\epsilon) \\ p_i(\epsilon) \end{bmatrix} = M_\epsilon \begin{bmatrix} q_i(0) \\ p_i(0) \end{bmatrix}, \ i = 1, \ldots, d. \tag{A.1}$$

Using the leapfrog scheme as an example, the matrix M_ϵ in (A.1) is

$$A = \begin{bmatrix} 1 - \epsilon^2/2 & \epsilon \\ -\epsilon + \epsilon^3/4 & 1 - \epsilon^2/2 \end{bmatrix}.$$

The chosen integration scheme is applied L times, so that the each element of the proposal obtains as

$$\begin{bmatrix} q_i^* \\ p_i^* \end{bmatrix} = M_\epsilon^L \begin{bmatrix} q_i(0) \\ p_i(0) \end{bmatrix}, \ i = 1, \ldots, d.$$

To obtain proposals that are independent of the current state, we consider combinations of ϵ and L so that the diagonal elements of M_ϵ^L are zero (i.e. $\epsilon L \approx \pi/2$), and get:

$$\begin{bmatrix} q_i^* \\ p_i^* \end{bmatrix} = \begin{bmatrix} 0 & r_L \\ s_L & 0 \end{bmatrix} \begin{bmatrix} q_i(0) \\ p_i(0) \end{bmatrix}$$

For the leapfrog method, the r_L and s_L are:

$$\begin{aligned} r_L &= 1 + \frac{1}{8}\epsilon^2 + \frac{3}{128}\epsilon^4 + \frac{5}{1024}\epsilon^6 + O(\epsilon^8), \\ s_L &= -1 + \frac{1}{8}\epsilon^2 + \frac{1}{128}\epsilon^4 + \frac{1}{1024}\epsilon^6 + O(\epsilon^8). \end{aligned}$$

Now, recall that the acceptance probability α:

$$\alpha = \min(1, \exp(\Delta)),$$

where $\Delta = (-H(q^*, p^*) + H(q, p))$. For dimension d, Δ simplifies as:

$$\Delta = \frac{1}{2}[\sum_{i=1}^{d}[(1 - s_L^2)q_i^2 + (1 - r_L^2)p_i^2]].$$

Note that from expressions for r_L and s_L, we have that $r_L \to 1$ and $s_L \to -1$ when $L \to \infty$. As a result of this, $\Delta \to 0$. This behaviour is in correspondence with what we expect, because it indicates that the acceptance probability will go towards one.

Owing to the fact that $q_i,\ p_i,\ i = 1, \ldots, d$ are iid standard normal, expressions for $\mu = \mathrm{E}(\Delta)$ and $\sigma^2 = \mathrm{Var}(\Delta)$ can be found for each of the integration schemes. Further, we use the central limit theorem so that

$$\frac{\Delta - \mu}{\sigma} \xrightarrow[d \to \infty]{} N(0, 1),$$

and this forms the basis for the formulas in Table 1.

References

Beskos, A., N. Pillai, G. Roberts, J.-M. Sanz-Serna, A. Stuart, et al. (2013). Optimal tuning of the hybrid monte carlo algorithm. *Bernoulli 19*(5A), 1501–1534.

Blanes, S., F. Casas, and J. Sanz-Serna (2014). Numerical integrators for the hybrid Monte Carlo method. *SIAM Journal on Scientific Computing 36*(4), A1556–A1580.

Carpenter, B., A. Gelman, M. Hoffman, D. Lee, B. Goodrich, M. Betancourt, M. A. Brubaker, J. Guo, P. Li, and A. Riddell (2015). Stan: a probabilistic programming language. *Journal of Statistical Software*.

Duane, S., A. D. Kennedy, B. J. Pendleton, and D. Roweth (1987). Hybrid Monte Carlo. *Physics letters B 195*(2), 216–222.

Girolami, M. and B. Calderhead (2011). Riemann manifold langevin and hamiltonian monte carlo methods. *Journal of the Royal Statistical Society: Series B (Statistical Methodology) 73*(2), 123–214.

Hoffman, M. D. and A. Gelman (2014). The No-U-Turn Sampler: adaptively setting path lengths in Hamiltonian Monte Carlo. *Journal of Machine Learning Research 15*(1), 1593–1623.

Lan, S., V. Stathopoulos, B. Shahbaba, and M. Girolami (2015). Markov chain monte carlo from lagrangian dynamics. *Journal of Computational and Graphical Statistics 24*(2), 357–378.

Leimkuhler, B. and S. Reich (2004). *Simulating Hamiltonian dynamics*, Volume 14. Cambridge University Press.

Liu, J. S. (2008). *Monte Carlo strategies in scientific computing*. Springer Science & Business Media.

Neal, R. M. (2010). MCMC using Hamiltonian dynamics. *Handbook of Markov Chain Monte Carlo 54*, 113–162.

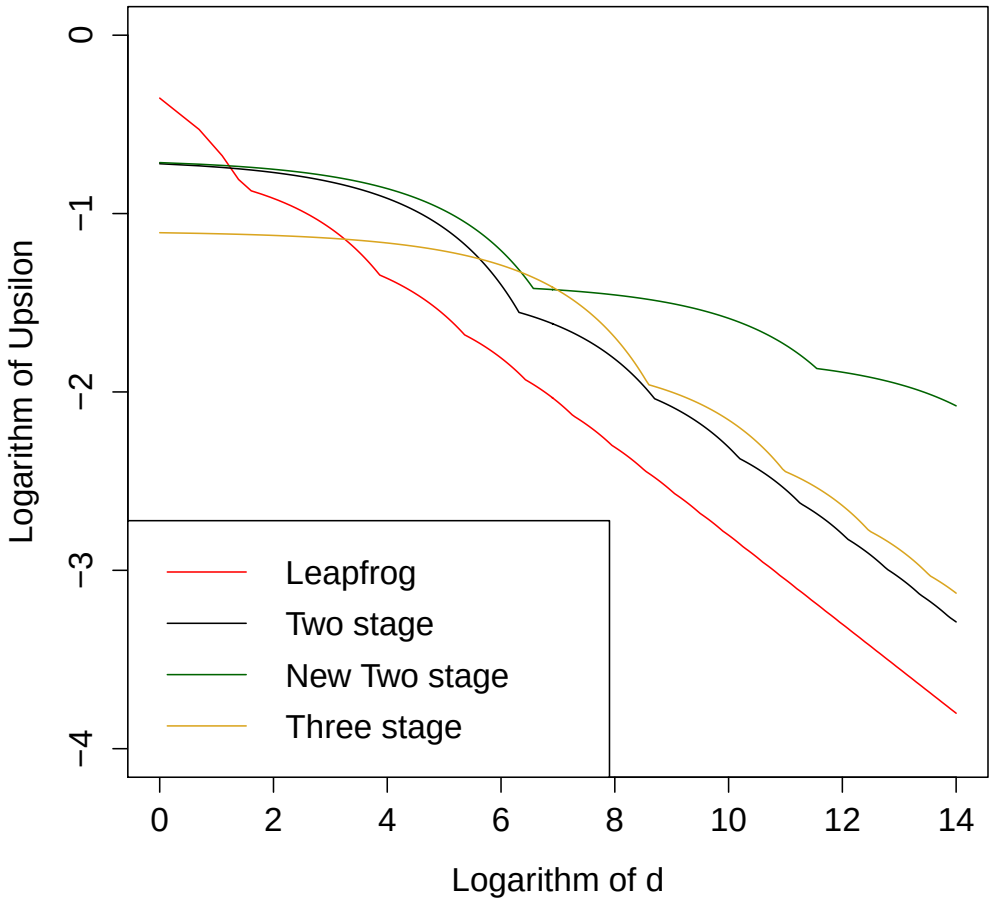

Figure 1: Logarithm of maximum expected number of accepted movements per calculation time as a function of $\log(d)$.

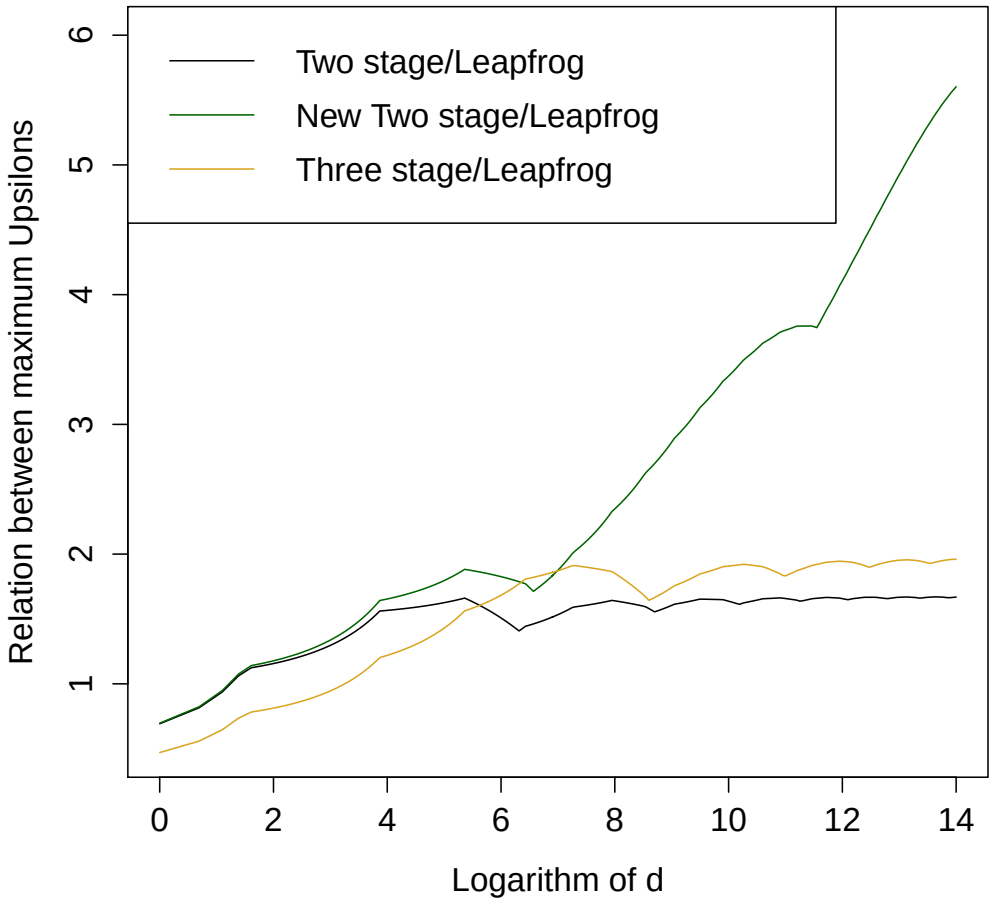

Figure 2: Relations between maximum expected number of accepted movements per calculation time as a function of $\log(d)$.

Table 1:

Table 1: Asymptotic Expected Acceptance Probability $\mathrm{E}(\alpha)$

Numerical scheme	$\mathrm{E}(\alpha)$	$\mathrm{E}_{\mathrm{New}}(\alpha)/\mathrm{E}_{\mathrm{Leapfrog}}(\alpha)$
Leapfrog	$2-2\Phi(\frac{1}{8}\epsilon^2\sqrt{d})$	1
Two-stage	$2-2\Phi(\frac{1}{24}(2-\sqrt{3})\epsilon^2\sqrt{d})$	$\approx 1+0.0908\epsilon^2\sqrt{d}+O(\epsilon^4)$
New two-stage	$2-2\Phi(\frac{1}{64}(5\sqrt{5}-11)\epsilon^4\sqrt{d})$	$\approx 1+0.0997\epsilon^2\sqrt{d}+O(\epsilon^4)$
Three-stage	$2-2\Phi(\frac{49087755}{19422491137}\sqrt{2}\epsilon^2\sqrt{d})$	$\approx 1+0.0969\epsilon^2\sqrt{d}+O(\epsilon^4)$

Table 2:

Numerical scheme	CPU time (s)	Integrator ϵ	minimum ESS	median ESS	maximum ESS	min ESS / CPU time	med ESS / CPU time
Australian credit data set ($d = 15$, $n = 690$)							
Leapfrog	15.45	0.0742	336	5000	5000	21.74	323.62
Two-stage	14.01	0.1911	465	5000	5000	33.19	**356.89**
New two-stage	14.50	0.1734	402	5000	5000	27.72	344.83
Three-stage	23.29	0.2979	1376	5000	5000	**59.08**	214.68
German credit data set ($d = 25$, $n = 1000$)							
Leapfrog	20.64	0.0386	2019	4755	5000	97.82	230.38
Two-stage	16.53	0.1110	2251	5000	5000	136.18	**302.48**
New two-stage	17.39	0.0999	2121	5000	5000	121.97	287.52
Three-stage	17.10	0.1866	4986	5000	5000	**291.58**	292.40
Heart data set ($d = 14$, $n = 270$)							
Leapfrog	7.63	0.1106	3364	4914	5000	440.89	644.04
Two-stage	6.46	0.2772	4950	5000	5000	**766.25**	733.99
New two-stage	6.22	0.2533	4369	5000	5000	702.41	**803.86**
Three-stage	7.86	0.4390	5000	5000	5000	636.13	636.13
Pima Indian data set ($d = 8$, $n = 532$)							
Leapfrog	12.52	0.0725	3687	4666	5000	294.49	372.68
Two-stage	8.95	0.1795	5000	5000	5000	**558.66**	**558.66**
New two-stage	9.07	0.1598	4365	4990	5000	481.26	550.17
Three-stage	9.38	0.2886	4986	5000	5000	531.56	508.65
Ripley data set ($d = 7$, $n = 250$)							
Leapfrog	18.18	0.1263	1351	1917	2793	74.31	105.45
Two-stage	17.37	0.2488	1513	2121	2967	87.10	122.12
New two-stage	18.50	0.2661	1547	2142	2787	83.62	115.78
Three-stage	16.55	0.3787	1765	2554	4144	**106.65**	**154.32**

Table 1: ESS and CPU times by numerical scheme for logistic regression applied to data sets from Girolami and Calderhead (2011). The values in bold face represent the best results within each data set.

Table 3:

Numerical scheme	CPU time (s)	Integrator ϵ	minimum ESS	median ESS	maximum ESS	min ESS / CPU time	med ESS / CPU time
			Dimension $d = 2$				
Leapfrog	13.80	0.7897	859	876	893	**62.25**	**63.48**
Two-stage	18.34	1.5863	962	969	977	52.45	52.84
New two-stage	20.41	1.3941	963	975	987	47.18	47.77
Three-stage	15.90	2.3257	955	963	971	60.06	60.57
			Dimension $d = 10$				
Leapfrog	48.59	0.4027	577	633	747	11.87	13.03
Two-stage	59.89	0.8027	605	645	748	10.10	10.77
New two-stage	69.54	0.7027	614	651	749	8.83	9.36
Three-stage	54.68	1.2399	696	731	845	**12.72**	**13.37**
			Dimension $d = 100$				
Leapfrog	853	0.2615	1180	1552	2416	1.38	1.82
Two-stage	1021	0.6171	1156	1533	2372	1.13	1.50
New two-stage	1162	0.5647	1242	1537	2353	1.07	1.32
Three-stage	699	1.0762	1228	1521	2333	**1.75**	**2.18**

Table 2: ESS and CPU times when using the four numerical schemes with a d-dimensional student t-model. The values in bold face represent the best results within each dimension.